

DETAILED ACTION

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

This application currently names joint inventors. In considering patentability of the claims under 35 U.S.C. 103(a), the examiner presumes that the subject matter of the various claims was commonly owned at the time any inventions covered therein were made absent any evidence to the contrary. Applicant is advised of the obligation under 37 CFR 1.56 to point out the inventor and invention dates of each claim that was not commonly owned at the time a later invention was made in order for the examiner to consider the applicability of 35 U.S.C. 103(c) and potential 35 U.S.C. 102(e), (f) or (g) prior art under 35 U.S.C. 103(a).

Claims 1, 10-12, 14-15, 18-20, 23-25, 29 and 30 are rejected under 35 U.S.C. 103(a) as being unpatentable over Ismael et al. (US 6,356,931) hereinafter Ismael, in view of E et al. (US 2004/ 0019639 A1) hereinafter E.

For claims 1 (method), 15 (apparatus), 20 (system) and 25 (an article of manufacture), Ismael teaches ***a computer-implemented method employed within a network of application server instances*** (e.g., Ismael's method of remote access to

management beans (see Abstract) is employed within an exemplary network of stations as illustrated in Fig. 1. Ismael mentions that Fig. 1 is a schematic representation of a multi-station network based system 1 with three stations, or nodes, or machines 3, 4 and 5 connected via a network 2. The network can have any desired structure. See c3:41-53. Furthermore, he says that Fig. 2 and 2A form a schematic representation of a computer server for a station of Fig. 1. See c3:5-6, and c3:67-c4:3. Therefore, it follows that the stations 3, 4 and 5 illustrated in Fig. 1 represents “*application server instances*” as claimed) **having a cluster architecture** (i.e., the term “cluster architecture” is not defined in the disclosure, but only an exemplary cluster architecture is illustrated. Therefore, without reading limitations from the specification into the claims, a broadest reasonable interpretation of the term “cluster architecture” in the context of a network is a plurality of computers inter-connected and grouped together in a network. Therefore, Fig. 1 of Ismael illustrates “a network of application server instances having a cluster architecture”), **comprising**:

displaying a representation of a plurality of management beans (MBeans) registered with an MBean server on a graphical user interface of a computing device, wherein each of the displayed MBeans represents a manageable resource of an application server instance within a cluster of application server instances

(Ismael teaches “a computer-implemented method for accessing from a client machine an object at a remote machine via a telecommunication network, the method comprising steps of:

- a) registering at least one object at the remote machine;

b) generating a machine page at the remote machine, which page contains at least one registered object; and

c) browsing the object via a network adaptor and the machine page at the remote machine using a browser at the client station.” Emphasis added, see c1:52-59.

Expanding further on the above method, Ismael further mentions, “The invention finds particular application to a network management system wherein the object to which access is sought is a managed object bean within a managed machine. The object can be one of a set of beans at the remote machine, whereby step (d) can comprise:

displaying at a client machine representations of beans at the remote machine which are modifiable remotely from the client machine;

responding to user selection at the client machine of a displayed bean representation to display at the client machine bean properties which are remotely modifiable; and

responding to user input at the client machine remotely to modify selected parameters of the bean.” Emphasis added, see c2:31-33.

Also referring to Fig. 3, he further mentions, “A managed object is a software abstraction of a resource that is controlled and

monitored by an agent. A managed object is referred to as a management bean or m-bean." Emphasis added, see c5:53-56.

Therefore, from the above excerpts it becomes clear that Ismael teaches ***displaying a representation of a plurality of management beans (MBeans)*** (e.g., see step (c) of the method above reciting "browsing the object..." and also step (d) of the method reciting "displaying at a client machine representations of beans at the remote machine". As highlighted in the above excerpts, the "object" referred to here is a managed object which is a management bean or m-bean) ***registered with an MBean server*** (e.g., see step (a) of the above method reciting "registering at least one object at the remote machine". Ismael further mentions, "A managed object in the agent is manageable as soon as it is registered with the framework." See c5:5) ***on a graphical user interface of a computing device*** (e.g. on a browser of a client station as mentioned in the excerpt above, "browsing the object via a network adaptor and the machine page at the remote machine using a browser at the client station." See c1:52-59), ***wherein each of the displayed MBeans represents a manageable resource of an application server instance within a cluster of application server instances*** (e.g., as shown in the excerpts above, Ismael mentions, "wherein the object to which access is sought is a managed object bean within a managed machine." Emphasis added, see c2:31-33. Also referring to Fig. 3, he further mentions, "A managed object is a software abstraction of a resource

that is controlled and monitored by an agent. A managed object is referred to as a management bean or m-bean." See c5:53-56.);

monitoring the management resources within the cluster, including receiving information regarding the manageable resources within the cluster from the plurality of MBeans registered with the MBean server (e.g., browsing the management bean objects representing manageable resources of the remote server instances of Fig. 1 from the client browser);

selecting one of the plurality of MBeans displayed in the graphical user interface (as mentioned in the excerpt above, "responding to user selection at the client machine of a displayed bean representation to display at the client machine bean properties which are remotely modifiable); ***and***

accessing an attribute of the selected MBean with the graphical user interface to view the received information regarding the manageable resource represented by the selected MBean (e.g., as mentioned in the excerpt above, "responding to user selection at the client machine of a displayed bean representation to display at the client machine bean properties which are remotely modifiable).

However, the claim further recites additional limitations ***each application server instance having a group of server nodes configured with a redundant set of application logic and associated data, the cluster having a dispatcher and a central service, the application server instances in communication with the***

central service, the central service having a locking service and a messaging service;

Ismael fails to teach employing his method of remotely monitoring manageable resources of application server instances using registered MBeans within a network of application server instances having the cluster architecture with all of the above additional limitations recited in the claim. In particular, Ismael fails to clearly address “*each application server instance having a group of server nodes*”, “*a redundant set of application logic and associated data*”, as well as “*a dispatcher*”, “*a central service*”, “*a locking service*” and “*a messaging service*” as recited in the claim. However, the Examiner notes that the instant specification does not provide any limiting definition for the terminology “*dispatcher*” utilized in the claim. Referring to Fig. 13, paragraph [00084] mentions, “*In one embodiment, dispatcher 1312 distributes service requests from clients to one or more of server nodes 1314, 1316, 1318 based on the load on each of the servers*”. Therefore, without limiting, one interpretation of the term “*dispatcher*” in light of one embodiment disclosed in the specification could be “*a module used for load balancing*”. The specification also does not provide any limiting definition for the terminology “*locking service*” and “*messaging service*” utilized in the claim. E teaches ***a cluster of application server instances*** (e.g., 104A and 104B in Fig. 1. E mentions, “*Enterprise server 102 may be implemented or clustered across one or more of the same computer systems as application servers 104, or on one or more separate computer systems*.” See [0035]), ***each application server instance***

within the cluster of application server instances having a group of server nodes configured with a redundant set of application logic and associated data (e.g., each application server instance 104 within the cluster of application server instances as illustrated in Fig. 1 has a group of processes 106, and associated data 108. These processes 106 can be interpreted as the claimed “server nodes” since they provide data and/or services for use by the clients. See [0035], wherein E mentions, “In one embodiment, the applications and/or processes within the application servers may provide data and/or services to enterprise server 102, for example, for use by the clients”. Or in the alternative, it would have been at least obvious to those of ordinary skill in the art to have some of these processes as server processes. Additionally, it is implicit or would have been at least obvious to those of ordinary skill in the art, according to the E reference that the processes 106 are, or can easily be, configured with a redundant set of application logic. This is because Fig. 1 illustrates a distributed data system, and regarding distributed data system, E explicitly mentions, “Distributed data systems may provide for load balancing and fail over to improve the overall quality of service of the system.” See [0006]. Thus it is implicit or at least would have been obvious to have the distributed data system as illustrated in Fig. 1 to provide load balancing and fail-over mechanism and thereby providing redundant set of logic in some of the processes 106), ***and a dispatcher*** (e.g., the module controlling the load balancing functionality) ***in communication with a central service*** (e.g., distributed store 110 in Fig. 1) ***associated with the cluster of***

application server instances, the central service having a locking service (e.g., lock mechanism 114 in Fig. 1) ***and a messaging service*** (see [0039] for messaging service for obtaining a token). Therefore, E teaches the above additional limitations of the claim not explicitly taught in Ismael. The question is whether it would have been obvious to one skilled in the art to combine the teaching of these two references. Ismael teaches a general technique for management of network resources using remote manipulation of management beans (Mbeans) similar to the technique claimed but does not use this technique for managing application server instances with the claimed clustered architecture, and E teaches a distributed data system having application server instances with the clustered network architecture as claimed but does not teach the technique for management of network resources using remote manipulation of management beans (Mbeans). However, the Examiner believes that since the general technique taught by Ismael for management of network resources is applicable and desirable for any system no matter what the purpose and/or architecture of the system, it would have been obvious to a person of ordinary skill in the art to employ the management technique of Ismael with the distributed data system of E so that the resources of the distributed data system can be monitored and managed efficiently, especially since Ismael explicitly suggests that his invention can be used in a network of any desired architecture (see *Ismael column 3 lines 51-53*).

Regarding claim 20, it is further noted that the claim has been interpreted as if 35 U.S.C 112, sixth paragraph, has been invoked.

For claims 10 and 11, Ismael implicitly teaches selecting one of the plurality of displayed MBeans with a pointing device or a keyboard (*9 in Fig. 2*).

Claims 12 (method), 18 (apparatus), 23 (system) and 29 (article of manufacture) are directed to accessing an attribute of an MBean representing a cluster manager of the network. However, the specification does not provide any limiting definition for the phrase "cluster manager" utilized in the claim. Referring to Fig. 12, paragraph [00077] mentions that "*each server node (e.g., 1218, 1228) includes...a cluster manager 1242, 1252 for communicating with messaging service 1204*". Therefore, without limiting, one interpretation of the term "cluster manager" in light of one embodiment disclosed in the specification is "a component or resource of the application servers used for communicating messages between the application servers and the central services of the system. As pointed out in the rejection of claim 1, E teaches that application servers can request locked access to a portion of primary data 112 from the distributed store 110 and the distributed store 110 can send a reply message to the application server including a token for the portion if the portion is not locked for another server (see [0039]). Therefore, E teaches a component or resource of the application servers used for communicating messages between the application servers and the central services of the system. Since Ismael teaches that all resources of a system worth monitoring can be represented as MBeans and since all MBeans can be represented and their attributes accessed using a graphical user interface, it would have been obvious to

those of ordinary skill in the art to combine the teaching of E and Ismael for accessing an attribute of an MBean representing a cluster manager of the network. The motivation for such combination would have been to monitor and manage lock requests within the distributed data system.

For claims 14 (method), 19 (apparatus), 24 (system) and 30 (article of manufacture), Ismael teaches invoking an operation of the selected MBean with the graphical user interface (*column 2 lines 29-30*).

Claim 13 is rejected under 35 U.S.C. 103(a) as being unpatentable over Ismael in view of E, further in view of Yeluripati et al. (US 7,086,065) hereinafter Yeluripati.

For claim 13, Ismael does not teach accessing a queue size attribute of the MBean representing the cluster manager to determine a number of requests waiting in the queue. However, using a queue to process requests is a well-known mechanism used in the art. E teaches that a request for lock may be queued by lock mechanism 114 (see *column 4 lines 2-3*). Yeluripati teaches a functional bean that receives client requests from a queue to service the request in a first come first serve basis (*column 7 lines 45-54*). Therefore, it would have been obvious to use a queue to service the requests in a MBean representing the cluster manager and subsequently access the queue size attribute of the MBean to determine a number of requests waiting in the queue. The motivation for using a queue would have been to serve the requests in a

first come first serve basis (*Yeluripati, column 7 lines 45-54*) and the motivation for accessing the queue size attribute would have been to monitor the cluster manager performance.

Claims 3-9, 17, 22, 27 and 28 are rejected under 35 U.S.C. 103(a) as being unpatentable over Ismael in view of E, further in view of Hessmer et al. (US2002/0112044) hereinafter Hessmer.

Ismael and E do not teach displaying a representation of a plurality of hierarchically organized MBeans as a tree structure having a root node, wherein the root node is an MBean representing the cluster of application server instances. they do not teach that the tree structure further includes one or more server nodes depending from the root node and showing kernel nodes, library nodes and service nodes depending from each of the one or more server nodes, wherein all these nodes are MBeans. Hessmer teaches a method and system for performing remote diagnostics on a process data access server, wherein he teaches displaying a set of diagnostic roots in the form of a hierarchical tree structure in the left pane of the graphical user interface associated with the diagnostic utility 100 (*Fig. 4, [0056]*). These diagnostic roots are elements to be monitored organized according to the type of elements. Hessmer's hierarchical tree structure organizes the presentation of the diagnostic roots having a root representing the cluster of servers and then showing a list of servers depending from the root and further showing various diagnostic roots depending from each of the servers. Therefore, it would have been obvious to a person of ordinary skill in the art to

incorporate this aspect of Hessmer's teaching with that of Ismael and E to represent the plurality of MBeans, each representing a manageable resource, in a hierarchical tree structure and organized in groups under respective server nodes as kernel, service and library nodes respectively based on the type of the resource. The motivation for using a hierarchical tree structure for representing the MBeans in various groups would have been to provide scalability of elements to expose lower levels and their associated information and further to provide ready access to a broad spectrum of diagnostic data via a graphical user interface (Hessmer, [0056]).

Response to Arguments

Applicants' arguments filed 05/21/2008 have been fully considered but they are not persuasive.

Applicants' have argued that nothing in Ismael discloses ***monitoring the manageable resources within the cluster of application server instances, including receiving information regarding the manageable resources within the cluster from the plurality of MBeans registered with the MBean server.***

The Examiner disagrees. Ismael teaches "a computer-implemented method for accessing from a client machine an object at a remote machine via a telecommunication network, the method comprising steps of:

- a) registering at least one object at the remote machine;

b) generating a machine page at the remote machine, which page contains at least one registered object; and

c) browsing the object via a network adaptor and the machine page at the remote machine using a browser at the client station." **See c1:52-59.**

Expanding further on the above method, Ismael further mentions, "The invention finds particular application to a network management system wherein the object to which access is sought is a managed object bean within a managed machine. The object can be one of a set of beans at the remote machine, whereby step (d) can comprise:

displaying at a client machine representations of beans at the remote machine which are modifiable remotely from the client machine;

responding to user selection at the client machine of a displayed bean representation to display at the client machine bean properties which are remotely modifiable; and

responding to user input at the client machine remotely to modify selected parameters of the bean." **Emphasis added, see c2:31-33.**

Also referring to Fig. 3, he further mentions, "A managed object is a software abstraction of a resource that is controlled and

monitored by an agent. A managed object is referred to as a management bean or m-bean." Emphasis added, see c5:53-56.

Therefore, from the above excerpts it becomes clear that Ismael teaches remotely monitoring manageable resources of application server instances using registered MBeans within a network of application server instances having a cluster architecture (note: although the cluster architecture is different than that recited in the claim, the Examiner did not rely on the particular cluster architecture of Ismael for the rejection. Instead it is the monitoring technique of resources using MBeans taught in the reference that was relied on). Thus the Examiner considers that monitoring manageable resources within a cluster of application server instances, including receiving information regarding the manageable resources within the cluster from plurality of MBeans registered with an MBean server would have been obvious based on the teaching of Ismael's disclosure.

Conclusion

THIS ACTION IS MADE FINAL. Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any

extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of this final action.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to RASHEDUL HASSAN whose telephone number is (571)272-9481. The examiner can normally be reached on M-F 7:30AM - 4PM EST.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Weilun Lo can be reached on 571-272-4847. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/Rashedul Hassan/
Examiner, Art Unit 2179

/Weilun Lo/
Supervisory Patent Examiner, Art Unit 2179